

# On the exactness of the cavity method for Weighted b-Matchings on Arbitrary Graphs and its Relation to Linear Programs

Mohsen Bayati\*    Christian Borgs\*    Jennifer Chayes\*    Riccardo Zecchina†

## Abstract

We consider the general problem of finding the minimum weight b-matching on arbitrary graphs. We prove that, whenever the linear programming relaxation of the problem has no fractional solutions, then the cavity or belief propagation equations converge to the correct solution both for synchronous and asynchronous updating.

Motivated by the cavity method, very fast distributed heuristic algorithms have recently been developed for the solution of random constraint satisfaction problems. In some cases, namely in the replica symmetric scheme, the algorithms generated by the cavity method are exactly of the form of a classic belief propagation (max-product or min-sum), i.e. a message-passing algorithm for efficiently computing marginal probabilities or finding the assignment with highest probability of a joint discrete probability distribution defined on a graph. The belief propagation (BP) algorithm converges to a correct solution if the associated graph is a tree, and may be also a good heuristic for some graphs with cycles. In other cases, e.g. when the space of solutions clusters into many subsets, the cavity method may lead to a more involved survey propagation (SP) algorithm [1, 2], in which some form of long range correlation among variables is included in the formalism.

In this paper, we study the problem of finding the minimum weight b-matchings in *arbitrary* graphs via the min-sum version of the cavity/ BP algorithm.

Let  $G = (V, E)$  be an undirected graph with edge weights  $w_{ij}$  for each edge  $\{i, j\} \in E$  and node capacities  $b_i$  for each node  $i \in V$ . The iterative message-passing algorithm based on synchronous BP for solving the weighted perfect b-matching problem is the following *simple* procedure: At each time, every vertex of the graph sends (real valued) messages to each of its neighbors. The message transmitted at time  $t$  from  $i$  to  $j$  is  $w_{ij}$  minus the  $b_i^{th}$  minimum of the messages previously received by  $i$  at time  $t - 1$  from all of its neighbors except  $j$ . At the end of each iteration, every vertex  $i$  selects  $b_i$  of its adjacent edges that correspond to the  $b_i$  smallest received messages. This procedure can be derived as the zero temperature limit of the cavity equations for the minimum b-matching problem. In what follows we will use as synonymous cavity equation, belief propagation (BP) or min sum equations.

We will show the following result: For arbitrary graphs  $G$ , and all sets of weights  $\{w_{ij}\}$ , after  $O(n)$  iterations, the set of selected edges converges to the correct solution, i.e., to the minimum weight perfect b-matching of  $G$ , provided that the Linear Programming (LP) relaxation of the problem has no fractional solutions. Additionally we introduce a new construction, a *generalized computation tree*, which allows us to analyze the more complicated case of BP with an asynchronous updating scheme, and prove convergence and correctness of it when each edge of the graph transmits at least  $\theta(n)$  messages.

Our proof gives new insight of the often-noted but poorly understood connection between the cavity method (or BP) and LP through the dual of the LP relaxation.

We also modify our BP algorithm and its analysis to include the problem of finding the *non-perfect* weighted b-matchings. Independent work on this aspect can be found in [3]. Previous exact results on *bipartite graphs* was obtained in [4] for  $b = 1$  and then was extended to all  $b$  in [5].

\*Microsoft Research; {mohsenb, borgs, jchayes}@microsoft.com

†Politecnico Di Torino; riccardo.zecchina@polito.it

The weighted b-matching problem is an important problem in combinatorial optimization. For extensive surveys see [6] and [7]. In physics, the study of the random 1-matching problem by the cavity method goes back to the work of Mèzard and Parisi [8] who made a celebrated conjecture for the expected optimum weight ( $\pi^2/6$ ) that was proven to be exact a decade later by Aldous [9]. BP algorithms have been the subject of extensive study in several communities (see [4] and [20] for a detailed survey of rigorous results about BP). Recent works have also suggested a connection between the BP algorithm and linear programming (LP) in particular problems. A relationship between iterative decoding of channel codes and LP decoding was studied in [10], [11], [12]. Other relationships were noted in the context of BP algorithms with convex free energies [13], [14], and in the case of BP algorithms for resource allocations [15]. For weighted 1-matchings, the connection was studied [16] in the context of similarities between BP equations and the primal-dual auction algorithm of Bertsekas [17]. Finally, we note that the BP equations for solving the weighted matching problem which we use in this paper have been previously studied in [16], [5]. These equations are also similar (though not identical) to the cavity equations for weighted matching problems and traveling salesman problems which are found in the statistical physics literature, see e.g. [8], [9], [18], and [19].

The main contributions of our results can be summarized as follows:

1. The properties of the BP equations for the weighted matching was first used in [4] and its correctness and convergence was shown for bipartite graphs with unique optimum solution. The same technique was used in [5] to extend the result to b-matchings in bipartite graphs. But this technique fails for graphs containing cycles with odd length. In order to bypass this difficulty we use a completely different tool, complementary slackness conditions of the LP relaxation and its dual, which is independent of the graph structure.
2. The connection between LP and the cavity method (or BP) is particularly important given that LP is a widely used technique for optimization. We show both convergence and correctness of the BP algorithm when LP relaxation has no fractional solutions, thus establishing a solid link between the two methods. In some recent work, [14], which generalizes methods of [13], the connection of the BP algorithm and LP relaxation is studied in the converged case of the BP. The authors also study interesting variations of the BP which have convex free energies.
3. *The asynchronous BP*, which includes the synchronous version as a special case, has been a more popular version for practical purposes. But, due to its more complicated structure, it has not been the subject of much rigorous study. Here we provide correctness and convergence proof of asynchronous BP for a combinatorial optimization problem based on the construction of a generalized computation tree, which can be used for the analysis of the both convergence and correctness of asynchronous message-passing algorithms as the cavity method.

Consider an undirected simple graph  $G = (V, E)$ , with vertices  $V = \{1, \dots, n\}$ , and edges  $E$ . Let each edge  $\{i, j\}$  have weight  $w_{ij} \in \mathbb{R}$ . Denote the set of neighbors of each vertex  $i$  in  $G$  by  $N(i)$ . Let  $\mathbf{b} = (b_1, \dots, b_n)$  be a sequence of positive integers such that  $b_i \leq \deg_G(i)$ . A subgraph  $M$  of  $G$  is called a *b-matching* (*perfect b-matching*) if the degree of each vertex  $i$  in  $M$  is at most  $b_i$  (equal to  $b_i$ ). Denote the set of b-matchings (perfect b-matchings) of  $G$  by  $M_G(\mathbf{b})$  ( $\text{PM}_G(\mathbf{b})$ ), and assume that it is non-empty. Clearly  $\text{PM}_G(\mathbf{b}) \subset M_G(\mathbf{b})$ .

The weight of a (perfect or non-perfect) b-matching  $M$ , denoted by  $W_M$ , is defined by  $W_M = \sum_{ij} w_{ij} 1_{\{i,j\} \in M}$ . In the next two sections, we will restrict ourselves to the case of *perfect* b-matchings. The analysis is extended to (possibly non-perfect) b-matchings in [20]. The minimum weight perfect b-Matching (b-MWPM),  $M^*$ , is defined by  $M^* = \arg\min_{M \in \text{PM}_G(\mathbf{b})} W_M$ . The goal of this paper is to find  $M^*$  via a min-sum belief propagation algorithm. Throughout the paper, we will assume that  $M^*$  is unique.

**Linear Programming Relaxation.** Assigning variables  $x_{ij} \in \{0, 1\}$  to the edges in  $E$ , we can express the weighted perfect b-matching problem as the problem of finding a vector  $\mathbf{x} \in \{0, 1\}^{|E|}$  that minimizes the total weight  $\sum_{ij \in E} x_{ij} w_{ij}$ , subject to the constraints  $\sum_{j \in N(i)} x_{ij} = b_i$  for all  $i \in V$ . Relaxing the constraint

that  $x_{ij}$  is integer, this leads to the following linear program and its dual:

$$\begin{array}{ccc|ccc}
\min & \sum_{\{i,j\} \in E} x_{ij} w_{ij} & & \max & \sum_{i=1}^n b_i y_i - \sum_{\{i,j\} \in E} \lambda_{ij} & \\
\text{subject to} & \sum_{j \in N(i)} x_{ij} = b_i & \forall i & \text{subject to} & w_{ij} + \lambda_{ij} \geq y_i + y_j & \forall \{i,j\} \in E \\
& 0 \leq x_{ij} \leq 1 & \forall \{i,j\} \in E & & \lambda_{ij} \geq 0 & \forall \{i,j\} \in E \\
& \text{Primal LP} & & & \text{Dual LP} & 
\end{array} \tag{1}$$

We say the LP relaxation (1) has *no fractional solution* if, every optimal solution  $x$  of LP satisfies  $x \in \{0, 1\}^{|E|}$ . Note that absence of fractional solutions implies uniqueness of integer solutions, since any convex combination of two integer solutions is a solution to the LP as well. We want to show that the BP algorithm for our problem converges to the correct solution, provided the LP relaxation (1) has no fractional solution.

**Complementary Slackness Conditions.** Complementary slackness for the LP and its dual state that the variables  $x^* = (x_{ij}^*)$  and  $y^* = (y_i^*)$ ,  $\lambda^* = (\lambda_{ij}^*)$  are optimum solutions to the LP relaxation and its dual (1), respectively, if and only if for all edges  $\{i, j\}$  of  $G$  both  $x_{ij}^*(w_{ij} + \lambda_{ij}^* - y_i^* - y_j^*) = 0$ . and  $(x_{ij}^* - 1)\lambda_{ij}^* = 0$  hold. Using the fact that the LP has no fractional solution, one can deduce the following modified complementary slackness conditions: For all  $\{i, j\} \in M^*$ ;  $w_{ij} + \lambda_{ij}^* = y_i^* + y_j^*$  and for all  $\{i, j\} \notin M^*$ ;  $\lambda_{ij}^* = 0$ .

By these conditions and the fact that  $\lambda_{ij}^* \geq 0$ , we have that  $w_{ij} \leq y_i^* + y_j^*$  for all  $\{i, j\} \in M^*$ , and  $w_{ij} \geq y_i^* + y_j^*$  for all  $\{i, j\} \notin M^*$ . However, as the counterexample given in [20] shows, it is in general not true that these inequalities are strict even when the LP has no fractional solution. Let  $S$  be the set of those edges in  $G$  for which  $|w_{ij} - y_i^* - y_j^*| > 0$ . We will assume the minimum gap is  $\epsilon$ . i.e.  $\epsilon = \min_{\{i,j\} \in S} |w_{ij} - y_i^* - y_j^*| > 0$ . Throughout this paper we assume that there exist an edge in  $G$  for which the strict inequality  $|w_{ij} - y_i^* - y_j^*| > 0$  holds and therefore  $\epsilon > 0$  is well defined. The other cases, where for each  $\{i, j\} \in E$  the equality  $w_{ij} = y_i^* + y_j^*$  holds, happens only for special cases are discussed in [20] and can be treated similarly. Let also

$$L = \max_{1 \leq i \leq n} |y_i^*|. \tag{2}$$

**Algorithm and Main Result.** The following algorithm is a synchronous implementation of BP for finding the minimum weight perfect b-matching (b-MWPM). The main intuition behind this algorithm (and, indeed, all BP algorithms) is that each vertex of the graph assumes the graph has no cycles, and makes the best (greedy) decision based on this assumption. Before applying the BP algorithm, we remove all *trivial* vertices from the graph. A vertex  $i$  is called trivial if  $\deg_G(i) = b_i$ . This is because all of the edges adjacent to  $i$  should be in every perfect b-matching. Therefore the graph can be simplified by removal of all trivial vertices and their adjacent edges.

---

#### Algorithm Sync-BP.

---

- (1) At times  $t = 0, 1, \dots$ , each vertex sends real-valued messages to each of its neighbors. The message of  $i$  to  $j$  at time  $t$  is denoted by  $m_{i \rightarrow j}(t)$ .
- (2) Messages are initialized by  $m_{i \rightarrow j}(0) = w_{ij}$  for all  $\{i, j\} \in E$  (the messages can actually be initialized by any arbitrary values [20])
- (3) For  $t \geq 1$ , messages in iteration  $t$  are obtained from messages in iteration  $t - 1$  recursively as follows:

$$\forall \{i, j\} \in E: \quad m_{i \rightarrow j}(t) = w_{ij} - b_i^{th\text{-min}}_{\ell \in N(i) \setminus \{j\}} \left[ m_{\ell \rightarrow i}(t-1) \right] \tag{3}$$

where  $k^{th\text{-min}}[A]$  denotes the  $k^{th}$  minimum<sup>1</sup> of set  $A$ .

---

<sup>1</sup>Note that the  $b_i^{th\text{-min}}_{\ell \in N(i) \setminus \{j\}}$  is well defined since we assumed that all trivial vertices are removed and thus there are at least  $b_i + 1$  elements in the set  $N(i)$  for each  $i$ .

- (4) The estimated b-MWPM at the end of iteration  $t$  is  $M(t) = \cup_{i=1}^n E_i(t)$  where  $E_i(t) = \{\{i, j_1\}, \dots, \{i, j_{b_i}\}\}$  is such that  $N(i) = \{j_1, j_2, \dots, j_{deg_G(i)}\}$  and  $m_{j_1 \rightarrow i}(t) \leq m_{j_2 \rightarrow i}(t) \leq \dots \leq m_{j_{deg_G(i)} \rightarrow i}(t)$ . i.e., among all  $i$ 's neighbors, choose edges to the  $b_i$  neighbors that transfer the smallest incoming messages to  $i$ .
- (5) Repeat (3)-(4) until  $M(t)$  converges<sup>2</sup>.

---

In Lemma 1, we will show the main intuition behind the equation (3) and how it is derived. But we note that one can also use the graphical model representations of [4], [5], [21] to obtain the standard BP equations for this problem, which, after some algebraic calculations, yield the recursive equation (3).

The main result of the paper says that the above algorithm, which is designed for graphs with no cycle (i.e., for trees), works correctly for a much larger family of graphs including those with many short cycles.

**Theorem 1** *Assume that the LP relaxation (1) has no fractional solution. Then the algorithm Sync-BP converges to  $M^*$  after at most  $\lceil \frac{2nL}{\epsilon} \rceil$  iterations.*

If the LP relaxation (1) has a fractional solution whose cost is strictly less than  $W_{M^*}$ , then [21], [3] have shown for the case of 1-matching that BP does not converge to  $M^*$ . It is straightforward to generalize this to perfect b-matching as well. But for the case in which the LP relaxation has a fractional solution whose cost is equal to  $W_{M^*}$ , BP fails in general. This is because the  $b_i^{th}$  minimum in equation (3) is not unique, and one needs an oracle to make the right decision. If such an oracle exists, then BP converges to  $M^*$ .

In what follows we first display the connection between the Sync-BP equations and the so-called computation tree. Next we discuss how the complementary slackness conditions are related to alternating paths in the graph  $G$ . These results are eventually used to prove that, when the LP relaxation has no fractional solutions, then solutions on the computation tree are the same as the solutions on the original graph  $G$ .

**Analysis of the Synchronous BP via Computation Tree.** The main idea behind the algorithm Sync-BP is that it assumes the graph  $G$  has no cycle. In other words, it finds the b-MWPM of a graph  $G'$  that has the same local structure as  $G$  but no cycles. The precise definition of the computational tree for Sync-BP goes as follows.

**Computation Tree.** For any  $i \in V$ , let  $T_i^t$  be the  $t$ -level computation tree corresponding to  $i$ , defined as follows:  $T_i^t$  is a weighted tree of height  $t + 1$ , rooted at  $i$ . All tree-nodes have labels from the set  $\{1, \dots, n\}$  according to the following recursive rules: (a) The root has label  $i$ , (b) The set of labels of the  $deg_G(i)$  children of the root is equal to  $N(i)$ , and (c) If  $s$  is a non-leaf node whose parent has label  $r$ , then the set of labels of its children is  $N(s) \setminus \{r\}$ .  $T_i^t$  is often called the *unwrapped tree* at node  $i$ . The computation tree is a well known technique for analyzing algorithms and constructed by replicating the local connectivity of the original graph. The messages received by node  $i$  in the belief propagation algorithm after  $t$  iterations in graph  $G$  are equivalent to those that would have been received by the root  $i$  in the computation tree, if the messages were passed up along the tree from the leaves to the root. A subtree  $\mathcal{M}$  of edges in the computation tree  $T_i^t$  is called a perfect *tree-b-matching* if for each *non-leaf* vertex with label  $i$  we have  $deg_{\mathcal{M}}(i) = b_i$ . Now denote the minimum weight perfect tree-b-matching (b-TMWPM) of the computation tree  $T_i^t$  by  $\mathcal{N}^*(T_i^t)$ . The following lemma shows that Sync-BP can be seen as a dynamic programming procedure that finds the minimum weight perfect tree-b-matching over the computation tree. Figure 1 shows a graph  $G$  and one of its corresponding computation tree.

**Lemma 1** *The algorithm Sync-BP solves the b-TMWPM problem on the computation tree. In particular, for each vertex  $i$  of  $G$ , the set of  $E_i(t)$  which was chosen at the end of iteration  $t$  by Sync-BP is exactly the set of  $b_i$  edges which are attached to the root in b-TMWPM of  $T_i^t$ .*

Lemma 1 characterizes the estimated b-MWPM,  $M(t)$ , and will be used later in the proof of the main result.

---

<sup>2</sup> The subgraph  $M(t)$  is not necessarily a perfect b-matching of  $G$  but we will show that after  $O(n)$  iterations it will be the minimum weight perfect b-matching.

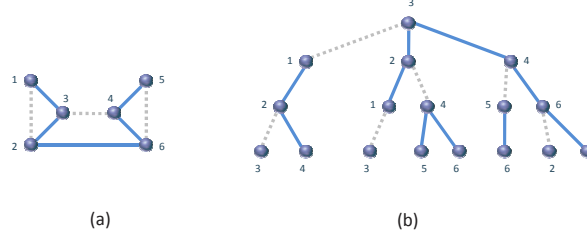


Figure 1: Part (a) shows a graph  $G$  where dashed and gray edges represent a 1-matching. Part (b) shows the computation tree  $T_3^2$  corresponding to  $G$  where the set of dashed and gray edges form a 1-TMWPM.

The following technical Lemma is also crucial for the proof of equivalence between BP and LP. It connects the complementary slackness conditions to paths on the graph  $G$  and on the computation tree. This lemma provides the connection between the absence of fractional solutions and the correctness of BP.

**Definition** A path  $P = (i_1, i_2, \dots, i_k)$  in  $G$  is called *alternating path* if it has the following two properties: (i) There exist a partition of edges of  $P$  into two sets  $A, B$  such that either  $(A \subset M^*, B \cap M^* = \emptyset)$  or  $(A \cap M^* = \emptyset, B \subset M^*)$ . Moreover  $A$  ( $B$ ) consists of all odd (even) edges; i.e.  $A = \{(i_1, i_2), (i_3, i_4), \dots\}$  ( $B = \{(i_2, i_3), (i_4, i_5), \dots\}$ ), and (ii) The path  $P$  might intersect itself or even repeat its own edges but no edge is repeated immediately. That is, for any  $1 \leq r \leq k-2$ :  $i_r \neq i_{r+1}$  and  $i_r \neq i_{r+2}$ .  $P$  is called an *alternating cycle* if  $i_1 = i_k$ .

**Lemma 2** Assume that the LP relaxation (1) has no fractional solution. Then for any alternating path  $P$  of length at least  $2n$ , there exists an edge  $\{i, j\} \in P$  such that the inequality  $|w_{ij} - y_i^* - y_j^*| > 0$  holds. That is,  $P \cap S \neq \emptyset$ .

**Proof of Theorem 1** We will prove Theorem 1, namely that if the LP relaxation (1) has no fractional solution and hence  $M^*$  is unique, then Sync-BP converges to the correct b-MWPM. We will do this by showing that if the depth of computation tree is large enough, then for any vertex  $i$ , its neighbors in  $M^*$  (b-MWPM of  $G$ ) are exactly those children that are selected in  $\mathcal{N}^*(T_i^t)$  (b-TMWPM of  $T_i^t$ ). Here is the main lemma that summarizes the above claim:

**Lemma 3** If the LP relaxation (1) has no fractional solution, then for any vertex  $i$  of  $G$  and for any  $t > \frac{2nL}{\epsilon}$ , the set of edges that are adjacent to root  $i$  in  $\mathcal{N}^*(T_i^t)$  are exactly those edges that are connected to  $i$  in  $M^*$ .

The proof of Lemma 3 is the main technical part of this work. The high level overview of the underlying argument goes as follows. Consider the computation tree  $(T_i^t)$  rooted at vertex  $i$  and look at  $\mathcal{N}^*(T_i^t)$ . We will assume that the claim of the lemma does not hold. That is, we assume that at the root,  $\mathcal{N}^*(T_i^t)$  does not choose the same edges as  $M^*$ -edges adjacent to  $i$ . Then we use the property of perfect tree-b-matchings, namely that each non-leaf vertex  $j$  is connected to exactly  $b_j$  of its neighbors, to construct a new perfect tree-b-matching on the computation tree. This new perfect tree-b-matching is going to have less total weight if the depth of the computation tree is large enough. This last step uses an alternating path argument which is a highly non-trivial generalization of the technique of [4] for the case of perfect 1-matching in bipartite graphs. For this part we will use the solutions to the dual LP (1).

**Proof of Lemma 3** Let us denote the lifting of a perfect b-matching  $M^*$  to a perfect tree-b-matching on  $T_i^t$  by  $\mathcal{M}^*$ . That is,  $\mathcal{M}^*$  consists of all edge of the computation tree with endpoint labels  $i, j$  such that  $\{i, j\} \in M^*$  as an edge in  $G$ . The goal is to show that  $\mathcal{N}^*(T_i^t)$  and  $\mathcal{M}^*$  have the same set of edges at the root of the computation tree. To lighten the notation, we denote the b-TMWPM of  $T_i^t$  by  $\mathcal{N}^*$ .

Assume the contrary, that there exist children  $i_{-1}, i_1$  of root  $i$  such that  $\{i, i_1\} \in \mathcal{M}^* \setminus \mathcal{N}^*$  and  $\{i, i_{-1}\} \in \mathcal{N}^* \setminus \mathcal{M}^*$ . Since both  $\mathcal{M}^*, \mathcal{N}^*$  are perfect tree-b-matchings, they have  $b_{i_1}$  edges connected to  $i_1$ . Therefore

there exist a child  $i_2$  of  $i_1$  such that  $\{i_1, i_2\} \in \mathcal{N}^* \setminus \mathcal{M}^*$ . Similarly there is a child  $i_{-2}$  of  $i_{-1}$  such that  $\{i_{-1}, i_{-2}\} \in \mathcal{M}^* \setminus \mathcal{N}^*$ . Therefore we can construct a set of alternating paths  $P_\ell$ ,  $\ell \geq 0$ , in the computation tree, that contain edges from  $\mathcal{M}^*$  and  $\mathcal{N}^*$  alternatively defined as follows. Let  $i_0 = \text{root } i$  and  $P_0 = (i_0)$  be a single vertex path. Let  $P_1 = (i_{-1}, i_0, i_1)$ ,  $P_2 = (i_{-2}, i_{-1}, i_0, i_1, i_2)$  and similarly for  $r \geq 1$ , define  $P_{2r+1}$  and  $P_{2r+2}$  recursively as follows:  $P_{2r+1} = (i_{-(2r+1)}, P_{2r}, i_{2r+1})$ ,  $P_{2r+2} = (i_{-(2r+2)}, P_{2r+1}, i_{2r+2})$  where  $i_{-(2r+1)}$ ,  $i_{2r+1}$  are nodes at level  $2r+1$  such that  $\{i_{2r}, i_{2r+1}\} \in \mathcal{M}^* \setminus \mathcal{N}^*$  and  $\{i_{-2r}, i_{-(2r+1)}\} \in \mathcal{N}^* \setminus \mathcal{M}^*$ . Similarly  $i_{-(2r+2)}$ ,  $i_{2r+2}$  are nodes at level  $2r+2$  such that  $\{i_{2r+1}, i_{2r+2}\} \in \mathcal{N}^* \setminus \mathcal{M}^*$  and  $\{i_{-(2r+1)}, i_{-(2r+2)}\} \in \mathcal{M}^* \setminus \mathcal{N}^*$ . Note that, by definition, such paths  $P_\ell$  for  $0 \leq \ell \leq t$  exist since the tree  $T_i^t$  has  $t+1$  levels and can support a path of length at most  $2t$  as defined above. Now consider the path  $P_t$  of length  $2t$ . It is an alternating path on the computation tree with edges from  $\mathcal{M}^*$  and  $\mathcal{N}^*$ . Let us refer to the edges of  $\mathcal{M}^*$  ( $\mathcal{N}^*$ ) as the  $\mathcal{M}^*$ -edges ( $\mathcal{N}^*$ -edges) of  $P_t$ . We will now modify the perfect tree-b-matching  $\mathcal{N}^*$  by replacing all  $\mathcal{N}^*$ -edges of  $P_t$  with their complement in  $P_t$  ( $\mathcal{M}^*$ -edges of  $P_t$ ). It is straightforward that this process produces a new perfect tree-b-matching  $\mathcal{N}'$  in  $T_i^t$ .

Let us assume, for the moment, the following lemma:

**Lemma 4** *The weight of the perfect tree-b-matching  $\mathcal{N}'$  is strictly less than that of  $\mathcal{N}^*$  on  $T_i^t$ .*

This completes the proof of Lemma 3 since Lemma 4 shows that  $\mathcal{N}^*$  is not the minimum weight perfect tree-b-matching on  $T_i^t$ , leading to a contradiction.  $\blacksquare$

**Proof of Lemma 4** It suffices to show that the total weight of the  $\mathcal{N}^*$ -edges of  $P_t$  is more than the total weight of  $\mathcal{M}^*$ -edges of  $P_t$ . For each vertex  $i_r \in P_t$  consider the value  $y_{i_r}^*$  from the optimum solution to the dual LP (1). Using the inequality  $w_{ij} \leq y_i^* + y_j^*$  for edges of  $\mathcal{M}^*$ , we obtain:  $\sum_{\{i,j\} \in P_t \cap \mathcal{M}^*} w_{ij} \leq \left( \sum_{r=-t}^t y_{i_r}^* \right) - y_{i_{(-1)t_t}}^* - k_1 \epsilon$  where  $k_1$  is the number of  $\mathcal{M}^*$ -edges of  $P_t$  that belong to  $S$ , i.e., the number of  $\mathcal{M}^*$ -edges of  $P_t$  endowed with the strict inequality  $w_{ij} \leq y_i^* + y_j^*$ , with a gap of at least  $\epsilon$ . On the other hand, using the inequality  $w_{ij} \geq y_i^* + y_j^*$  for edges of  $\mathcal{N}^*$  we have:  $\sum_{\{i,j\} \in P_t \cap \mathcal{N}^*} w_{ij} \geq \left( \sum_{r=-t}^t y_{i_r}^* \right) - y_{i_{(-1)t+1_t}}^* + k_2 \epsilon$  where now  $k_2$  is number of  $\mathcal{N}^*$ -edges of  $P_t$  that belong to  $S$ , or equivalently the number of times the inequality  $w_{ij} \geq y_i^* + y_j^*$  is strict with a gap of at least  $\epsilon$ . One finds

$$\sum_{\{i,j\} \in P_t \cap \mathcal{N}^*} w_{ij} - \sum_{\{i,j\} \in P_t \cap \mathcal{M}^*} w_{ij} = y_{i_{(-1)t_t}}^* - y_{i_{(-1)t+1_t}}^* + (k_1 + k_2) \epsilon \stackrel{(a)}{\geq} (k_1 + k_2) \epsilon - 2L \stackrel{(b)}{\geq} (k_1 + k_2) \epsilon - 2L \stackrel{(c)}{>} 0.$$

Here (a) uses definition of  $L$  from eq. 2 and (b) uses the fact that for all  $i, j$ :  $\lambda_{ij}^* \geq 0$ . The main step is (c), which uses Lemma 2 as follows. Path  $P_t$  has length  $2t$ , and each continuous piece of it with length  $2n$  has a projection to the graph  $G$  which satisfies the conditions of Lemma 2. This means the path has at least one edge from the set  $S$ . Thus  $(k_1 + k_2) \geq \frac{2t}{2n} > \frac{2L}{\epsilon}$ . This completes the proof of Lemma 4.  $\blacksquare$

**Extension to Possibly Non-Perfect b-Matchings.** Here we note that the algorithm and the results of the previous sections can be easily generalized to the case of b-matchings (subgraphs  $H$  of  $G$  such that degree of each vertex  $i$  in  $H$  is at most  $b_i$ ). Let  $U(H) \subset V$  be the set of *unsaturated* vertices of  $G$  (vertices  $i \in V$  such that  $\deg_H(i) < b_i$ ). The minimum weight b-Matching (b-MWM),  $H^*$ , is the b-Matching such that  $H^* = \arg\max_{H \in \mathcal{M}_G(b)} W_H$ . Note that  $H^*$  does not include any edge with positive weight because removing such edges from  $H^*$  reduces its weight while keeping it a b-matching.

**Asynchronous BP.** In the remaining we study the asynchronous version of the BP algorithm. The update equations are exactly analogous to the synchronous version, but at each time only a subset of the edges are updated in an arbitrary order. Consider the set  $\vec{E}$  of all directed edges in the  $G$ ; i.e.,  $\vec{E} = \{(i \rightarrow j) \text{ s.t. } i \neq j \in V\}$ . Let  $A$  be a sequence  $\vec{E}(1), \vec{E}(2), \dots$  of subsets of the set  $\vec{E}$ . Then the asynchronous BP algorithm corresponding to the sequence  $A$  can be obtained by modifying only the update rule in the step (3) of the algorithm Sync-BP as follows:  $m_{i \rightarrow j}(t) = w_{ij} - b_i^{th\text{-min}}_{\ell \in N(i) \setminus \{j\}} \left[ m_{\ell \rightarrow i}(t-1) \right]$  if  $(i \rightarrow j) \in \vec{E}(t)$  and if  $(i \rightarrow j) \notin \vec{E}(t)$  then the message will not be updated, i.e. it remains equal to  $m_{i \rightarrow j}(t-1)$ .

This is the most general form of the asynchronous BP and it includes the synchronous version ( $\vec{E}(t) = \vec{E}$  for all  $t = 1, 2, \dots$ ) as a special case. In many applications, a special case of the asynchronous BP is used for which each set  $\vec{E}(t)$  consists of a single element.

We assume that the sequence  $A$  of the updates does not have *redundancies*. That is, no edge direction  $(i \rightarrow j) \in \vec{E}$  is re-updated before at least one of its incoming edge directions  $((\ell \rightarrow i) \text{ for } \ell \in N(i) \setminus \{j\})$  is updated. More formally, if  $(i \rightarrow j) \in \vec{E}(t) \cap \vec{E}(t+s)$  and  $(i \rightarrow j) \notin \cup_{r=1}^{s-1} \vec{E}(t+r)$ , then at least for one  $\ell \in N(i) \setminus \{j\}$ , we should have  $(\ell \rightarrow i) \in \cup_{r=1}^{s-1} \vec{E}(t+r)$ .

Let us denote the above algorithm by Async-BP. We claim that, if each edge direction  $(i \rightarrow j) \in \vec{E}$  is updated  $\theta(n)$  times, then the same result as Theorem 1 can be proved here. That is, let  $u(t)$  be the minimum number of times that an edge direction of the graph  $G$  appears in the sequence  $\vec{E}(1), \dots, \vec{E}(t)$ ; i.e.,  $u(t) = \min_{(i \rightarrow j) \in \vec{E}} \left( \left| \{ \ell : s.t. \ 1 \leq \ell \leq t \text{ and } (i \rightarrow j) \in \vec{E}(\ell) \} \right| \right)$ . From the definition,  $u(t)$  is a non-decreasing function of  $t$ . We claim that the following result holds:

**Theorem 2** *Assume that the LP relaxation (1) has no fractional solution. Then the algorithm Async-BP converges to  $M^*$  after at most  $t$  iterations, provided  $u(t) > \frac{2nL}{\epsilon}$ .*

Proof of the above theorem relies on the notion of generalized computation tree for the asynchronous version of the BP algorithm which will be given in the longer version of this paper [20].

Finally we note that the same algorithm as Async-BP and the same result as Theorem 2 can be stated and proved for the (possibly non-perfect) b-matchings as well.

## 1 Acknowledgements

We would like to thank László Lovász, Andrea Montanari, Elchannan Mossel and Amin Saberi for useful discussions. This work was done while Riccardo Zecchina was a Visiting Researcher in the Theory Group at Microsoft Research, and was supported by the Microsoft Technical Computing Initiative.

## References

- [1] M. Mezard and R. Zecchina “Random K-satisfiability: from an analytic solution to a new efficient algorithm,” *Phys.Rev. E E*, 66, 056126, 2002.
- [2] A. Braunstein, M. Mezard, and R. Zecchina, “Survey propagation: an algorithm for satisfiability,” *Random Structures and Algorithms*, vol. 27, pp. 201–226, 2005.
- [3] S. Sanghavi, D. Malioutov, A. Willsky “Linear programming analysis of loopy belief propagation for weighted matching”, to appear in *NIPS*, 2007.
- [4] M. Bayati, D. Shah, and M. Sharma, “Maximum weight matching via max-product belief propagation,” *Preliminary version appeared at IEEE ISIT 2005. Longer version to appear in IEEE Trans. Information Theory*, 2007.
- [5] B. Huang, T. Jebara, “Loopy belief propagation for bipartite maximum weight b-matching”, *Artificial Intelligence and Statistics (AISTATS)*, March, 2007.
- [6] A. Gerards, “Matching. Volume 7 of ,” *Hand book of Operation Research and Management Science*, Chapter 3, pp. 135-224. North-Holland, 1995.
- [7] W. Pulleyblank, “Matchings and extensions.” Volume 1 of *Handbook of Combinatorics*, Chapter 3, pp. 179-232, North Holland, 1995.
- [8] M. Mezard and G. Parisi “Mean-field equations for the matching and travelling Salesman problems,” *Eurhophysics letters*, Vol. 2, pp. 913-918, 1986.
- [9] D. Aldous, “The zeta (2) Limit in the Random Assignment Problem,” *Random Structures and Algorithms*, Vol. 18, pp. 381-418, 2001.

- [10] J. Feldman, M. Wainwright and D. Karger, “Using linear programming to decode binary linear codes”, *IEEE Transactions on Information Theory*, vol. 51, pp. 954-972, 2005.
- [11] P.O. Vontobel and R. Koetter, “On the relationship between linear programming decoding and min-sum algorithm decoding,” *Proc. ISITA 2004*, Parma, Italy, pp. 991–996, Oct. 10-13, 2004.
- [12] P.O. Vontobel and R. Koetter, “Graph-cover decoding and finite-length analysis of message-passing iterative decoding of LDPC codes,” to appear in *IEEE Trans. Inform. Theory*, <http://www.arxiv.org/abs/cs.IT/0512078>.
- [13] M. Wainwright, T. Jaakkola, and A. Willsky, “MAP estimation via agreement on trees: message-passing and linear programming”, *IEEE Transactions on Information Theory*, 51(11):3697:3711, 2005.
- [14] Y. Weiss, C. Yanover and T. Meltzer “MAP Estimation, Linear Programming and Belief Propagation with Convex Free Energies,” *UAI*, 2007.
- [15] C. Moallemi and B. Van Roy, “A Message-Passing Paradigm for Resource Allocation,” preprint June 2007.
- [16] M. Bayati, D. Shah, and M. Sharma, “Max-product for maximum weight matching: convergence, correctness and LP duality,” in *IEEE Int. Symp. Information Theory*, 2006.
- [17] D. P. Bertsekas, “The auction algorithm: A distributed relaxation method for the assignment problem,” *Annals of Operations Research*, vol. 14, 1988.
- [18] D. Gamarnik, T. Nowicki and G. Swirszcz, “Maximum Weight Independent Sets and Matchings in Sparse Random Graphs. Exact Results using the Local Weak Convergence Method”, *Random Structures and Algorithm*, Vol.28, No. 1, pp. 76-106, 2005.
- [19] L. Zdeborová and M. Mézard, “The number of matchings in random graphs”, *J. Stat. Mech.*, 2006.
- [20] M. Bayati, C. Borgs, J. Chayes and R. Zecchina, “Belief-Propagation for Weighted b-Matchings on Arbitrary Graphs and its Relation to Linear Programs with Integer Solutions”, in *arXiv*, <http://www.arxiv.org/abs/0709.1190v2>, September 8, 2007.
- [21] S. Sanghavi, “Equivalence of LP Relaxation and Max-Product for Weighted Matching in General Graphs”, *IEEE Information Theory Workshop*, September 2007, preliminary version arXiv:0705.0760, May 5 2007.